

**Best
Available
Copy**

ABSTRACT

In this paper, we apply the Perron-Frobenius theory for non-negative matrices to the analysis of variance asymptotics for simulations of finite state Markov chains to which importance sampling is applied. The results show that we can typically expect the variance to grow (at least) exponentially rapidly in the length of the time horizon simulated. The exponential rate constant is determined by the Perron-Frobenius eigenvalue of a certain matrix. Applications to cumulative costs, terminal costs, steady-state costs, and the likelihood ratio gradient estimator are presented. In addition, the implications for general discrete-event simulations are presented.

KEYWORDS:

importance sampling, Markov chains, discrete-event simulation

IMPORTANCE SAMPLING FOR MARKOV CHAINS:
ASYMPTOTICS FOR THE VARIANCE

by

Peter W. Glynn

TECHNICAL REPORT No. 73

OCTOBER 1991

Prepared under the Auspices
of
U.S. Army Research Contract
DAAL-03-91-G-0101

Accession Number	
DTIC Number	
DTIC 123	
Availability Codes	
Distribution/	
Availability Codes	
Avail and/or	
Dist	Special
A-1	



Approved for public release: distribution unlimited.

Reproduction in whole or in part is permitted for any
purpose of the United States government.

DEPARTMENT OF OPERATIONS RESEARCH
STANFORD UNIVERSITY
STANFORD, CALIFORNIA

92-11647



92 4 28 162

ABSTRACT

In this paper, we apply the Perron-Frobenius theory for non-negative matrices to the analysis of variance asymptotics for simulations of finite state Markov chains to which importance sampling is applied. The results show that we can typically expect the variance to grow (at least) exponentially rapidly in the length of the time horizon simulated. The exponential rate constant is determined by the Perron-Frobenius eigenvalue of a certain matrix. Applications to cumulative costs, terminal costs, steady-state costs, and the likelihood ratio gradient estimator are presented. In addition, the implications for general discrete-event simulations are presented.

KEYWORDS:

importance sampling, Markov chains, discrete-event simulation

Research supported by the U.S. Army Research Office under Contract DAAL-03-91-G-0101, National Science Foundation Grant DDM-9101580 and the Stanford Institute for Manufacturing Automation.

1. INTRODUCTION

Suppose that we are given a real-valued finite-horizon performance measure $Y = f(X_0, X_1, \dots, X_n)$, where $X = (X_n : n \geq 0)$ is a Markov chain having initial distribution μ and transition matrix P . (We assume throughout this paper that the state space S is finite, unless otherwise stated.) The standard way to estimate the expected performance $\alpha = EY$ is to use a sample mean of i.i.d. replicates of the r.v. Y generated under initial distribution μ and transition matrix P .

However, importance sampling offers an alternative. Specifically, let ν and K be, respectively, an initial distribution and transition matrix chosen so that $\mu(x) > 0$ implies $\nu(x) > 0$ and $P(x, y) > 0$ implies $K(x, y) > 0$. To indicate the dependence of the expectation operator of the Markov chain X upon the initial distribution and transition matrix, we shall write $E_P(\cdot)$ to denote the expectation operator relative to initial distribution μ and transition matrix P , whereas $E_K(\cdot)$ is the expectation operator in which the initial distribution and transition matrix are given by ν and K , respectively. Then, it is easily seen that $\alpha = E_P Y$ can be written as

$$\begin{aligned}
 \alpha &= \sum_{x_0, \dots, x_n} f(x_0, \dots, x_n) \mu(x_0) \prod_{i=0}^{n-1} P(x_i, x_{i+1}) \\
 (1.1) \quad &= \sum_{x_0, \dots, x_n} \left[f(x_0, \dots, x_n) \frac{(\mu x_0)}{\nu(x_0)} \prod_{i=0}^{n-1} \frac{P(x_i, x_{i+1})}{K(x_i, x_{i+1})} \right] \cdot \nu(x_0) \prod_{i=0}^{n-1} K(x_i, x_{i+1}) \\
 &= E_K Y L_n
 \end{aligned}$$

where

$$L_n = \frac{\mu(X_0)}{\nu(X_0)} \prod_{i=0}^{n-1} \frac{P(X_i, X_{i+1})}{K(X_i, X_{i+1})}.$$

The r.v. L_n is known as the likelihood ratio (of $E_P(\cdot)$ relative to $E_K(\cdot)$). The identity (1.1) suggests that α can be estimated by calculating i.i.d. replicates of the r.v. $Y L_n$ generated under initial distribution ν and transition matrix K . This technique is known as the method of (static) importance sampling for Markov chains. (See Glynn and Iglehart (1989) for a discussion of a related variant of importance sampling known as dynamic importance sampling.)

Our goal in this paper is to show that this importance sampling technique is always very poorly behaved from a variance standpoint when the time horizon is large. In fact,

we will show that for a wide variety of performance measures, the variance essentially grows exponentially rapidly in the length n of the time horizon. This result suggests that (static) importance sampling of the kind described above is typically going to be (highly) inefficient when the time horizon is large. While this conclusion is suggested by some previous analyses (see, for example, Glynn (1987)), our current treatment permits a rather precise quantitative characterization of the statistical inefficiency of importance sampling in such a context.

These results are of interest in at least three different problem settings.

Setting 1 (Variance Reduction): By choosing ν and K judiciously, one hopes that the estimator obtained via importance sampling will have a corresponding variance that is significantly lower than that of the standard estimator. This expectation is borne out in certain applications contexts. (See, for example, Goyal et al (1990).) However, our results suggest that one needs to proceed with caution in applying importance sampling to performance measures in which the time horizon is large.

Setting 2 (Gradient Estimation): As is well known (see, for example, Glynn (1986a) and Glynn (1987)), formula (1.1) underlies the likelihood ratio method for estimating gradients of mean performance measures with respect to vectors of continuous decision parameters. In Section 6 of this paper, we will discuss the implications of our results for such gradient estimation problems. Our results will show that when the time horizon corresponding to the performance measure is large, there is a best possible choice for K and we will identify the optimal transition matrix.

Setting 3 (Optimization): The likelihood ratio gradient estimator mentioned above can be used for optimization purposes. Two different approaches can be followed. One method involves developing a stochastic recursive algorithm (such as the Robbins-Monro stochastic approximation scheme) which is driven by newly generated likelihood ratio gradient estimators at each iteration (see, for example Glynn (1986b)). The second idea is to use importance sampling to generate via simulation, at a single point in the decision parameter space, an estimate to the entire response surface. Typically, the estimated response surface

will be (at least) twice continuously differentiable. Thus, one can attempt to estimate the optimizer of the true surface by applying deterministic optimization algorithms (such as Newton's method) to the estimated response surface (see Rubinstein and Shapiro (1989) for further details). It turns out that the gradients associated with the estimated response surface are precisely the likelihood ratio gradient estimators of Setting 2. In Section 6, we describe this connection more fully and discuss the implications of our results for this approach.

The remainder of this paper is organized as follows. In Section 2, we describe a formula for the variance of an arbitrary estimator obtained via importance sampling. This formula is fundamental to our subsequent analysis. In Sections 3 through 6, we apply the formula to: cumulative costs (Section 3), terminal costs (Section 4), steady-state costs (Section 5), and likelihood ratio gradient estimators (Section 6). Finally, in Section 7, we describe the implications of the theory developed in this paper for general discrete-event simulations.

2. A FORMULA FOR THE VARIANCE

Given that we apply importance sampling with initial distribution ν and transition matrix K to the estimation of α , the variance will be given by

$$(2.1) \quad \begin{aligned} \text{var}_K[YL_n] &= E_K Y^2 L_n^2 - (E_K Y L_n)^2 \\ &= E_K Y^2 L_n^2 - \alpha^2. \end{aligned}$$

Since α is independent of K , our goal is to simplify the expression for $E_K Y^2 L_n^2$. In pursuit of this objective, we note that

$$(2.2) \quad E_K Y^2 L_n^2 = \sum_{x_0, \dots, x_n} f^2(x_0, \dots, x_n) \frac{\mu^2(x_0)}{\nu(x_0)} \prod_{i=0}^{n-2} \frac{P^2(x_i, x_{i+1})}{K(x_i, x_{i+1})}$$

Let $G = (G(x, y) : x, y \in S)$ be the matrix in which $G(x, y) = P^2(x, y)/K(x, y)$ when $P(x, y) > 0$ and $G(x, y) = 0$ when $P(x, y) = 0$. Assume that the stochastic matrix P is irreducible. Then, G is necessarily irreducible. Since G is clearly non-negative, we may apply the Perron-Frobenius theory for non-negative matrices to the study of G . In particular (see Karlin and Taylor (1975)), it can be asserted that G possesses a positive eigenvalue $\lambda = \lambda(G)$ (known as the Perron-Frobenius eigenvalue) such that λ is the eigenvalue of

maximum (complex) modulus. Furthermore, there exists precisely one linearly independent column eigenvector associated with λ . This eigenvector h can be chosen to be strictly positive in all components. Since h is the eigenvector corresponding to λ , it follows that

$$\sum_y G(x, y)h(y) = \lambda h(x)$$

for $x \in S$. Hence,

$$\sum_y G(x, y) \frac{h(y)}{\lambda h(x)} = 1.$$

Let $R = (R(x, y) : x, y \in S)$ be the matrix in which $R(x, y) = G(x, y)h(y)/(\lambda h(x))$. Then, R is non-negative with row sums equal to 1, and is hence stochastic. Furthermore, R is irreducible since G is. Noting that

$$G(x, y) = \lambda h(x)R(x, y)/h(y),$$

it is evident that

$$(2.3) \quad \prod_{i=0}^{n-1} G(x_i, x_{i+1}) = \lambda^n \frac{h(x_0)}{h(x_n)} \prod_{i=0}^{n-1} R(x_i, x_{i+1}).$$

Let $\eta = (\eta(x) : x \in S)$ be the stochastic vector defined by

$$\eta(x) = \gamma^{-1} \mu^2(x)/\nu(x),$$

where $\gamma = \sum_x \mu^2(x)/\nu(x)$. With the aid of (2.3), we may now express (2.2) as

$$\begin{aligned} E_K Y^2 L_n^2 &= \gamma \lambda^n \sum_{x_0, \dots, x_n} f^2(x_0, \dots, x_n) \eta(x_0) \prod_{i=0}^{n-1} R(x_i, x_{i+1}) h(x_0)/h(x_n) \\ &= \gamma \lambda^n E_R \left[f^2(X_0, \dots, X_n) \frac{h(X_0)}{h(X_n)} \right], \end{aligned}$$

where $E_R(\cdot)$ is the expectation operator associated with initial distribution η and transition matrix R . We summarize our discussion with the following variance identity; the specialization of this formula to r.v.'s Y that are additive functionals is implicit in much of the large deviations discussion given in Bucklew (1990).

THEOREM 1. If P is irreducible, then

$$\text{var}_K[YL_n] = \gamma\lambda^n E_R[Y^2 h(X_0)/h(X_n)] - \alpha^2,$$

where γ, λ, h , and R are defined as above.

Typically, the magnitude of Y is polynomial in n (see Sections 3 through 6). Thus, the variance of YL_n is determined by the exponential behavior of λ^n (if $\lambda \neq 1$). Our next proposition tells us that λ is typically strictly greater than 1. Hence, if Y is of polynomial order in n , it is evident that $\text{var}_K[YL_n]$ is basically increasing geometrically fast at rate λ .

PROPOSITION 1. The quantities γ and λ are always greater than or equal to 1. Also, $\gamma > 1$ if $\mu \neq \nu$. Furthermore, suppose P is irreducible. Then, $\lambda > 1$ if $P \neq K$.

For the proof of Proposition 1, see the Appendix. Thus, we may conclude that in any non-trivial importance sampling context (i.e. $P \neq K$), the sequence of multipliers $\gamma\lambda^n$ is growing geometrically fast. Hence, in order that $\text{var}_K[YL_n]$ be well-behaved as a function of n , it is evident that $E_R[Y^2 h(X_0)/h(X_n)]$ must be small. For example, if $Y = I(A)$, where A is a "rare event" under R , $\text{var}_K[YL_n]$ can still be of moderate size (see Cottrell et al. (1983)). However, as we will see in Sections 3 through 6, importance sampling in most other problem settings leads to geometric growth in the variance.

3. APPLICATION 1: CUMULATIVE COSTS

Let f be a real-valued function defined on the state space S of the Markov chain X . Suppose that the performance measure is the cumulative cost corresponding to the function f , namely

$$Y_n = \sum_{k=0}^{n-1} f(X_k).$$

Assume P is irreducible and aperiodic. Then, Theorem 1 applies that

$$(3.1) \quad \text{var}_K[YL_n] = \gamma\lambda^n E_R[Y_n^2 h(X_0)/h(X_n)] - \alpha_n^2$$

where $\alpha_n = E_P Y_n$. Since $|Y_n| \leq n \cdot \|f\|$, where $\|f\| = \max\{|f(x)| : x \in S\}$, it is evident that $\alpha_n^2 = O(n^2)$ as $n \rightarrow \infty$. If $K \neq P$, Proposition 1 states that $\lambda > 1$ so that it is evident that the growth of the variance is governed by $\gamma\lambda^n E_R[Y_n^2 h(X_0)/h(X_n)]$.

Note that R inherits the irreducibility and aperiodicity of P . Then, we can assert that R has a unique stationary distribution $\pi_R(\cdot)$. Let $P_R(\cdot) = E_R[I(\cdot)]$ be the probability distribution on X associated with initial distribution η and transition matrix R . We can apply Theorem 4 of Niemi and Nummelin (1982) to conclude that there exist finite constants $\beta = \sum_x \pi_R(x)f(x)$ and σ_R^2 such that for any $t \in R$ and $x, y \in S$,

$$\begin{aligned} P_R\{n^{\frac{1}{2}}(Y_n/n - \beta) \leq t, X_n = z | X_0 = x\} \\ \rightarrow P\{N(0, \sigma_R^2) \leq t\} \pi_R(z) \end{aligned}$$

as $n \rightarrow \infty$, where $N(0, \sigma_R^2)$ is a normal r.v. with mean zero and variance σ_R^2 . Since (3.2) holds for all x , it is evident that

$$\begin{aligned} P_R\{X_0 = x, n^{\frac{1}{2}}(Y_n/n - \beta) \leq t, X_n = z\} \\ \rightarrow \eta(x)P\{N(0, \sigma_R^2) \leq t\} \pi_R(z) \end{aligned}$$

as $n \rightarrow \infty$. Note that the r.v.'s X_0 , $n^{\frac{1}{2}}(Y_n/n - \beta)$, and X_n are all asymptotically independent under $P_R(\cdot)$. The continuous mapping principle, together with a converging-together argument (see, for example, Billingsley (1968)), therefore proves that under the distribution $P_R(\cdot)$,

$$(3.4) \quad (Y_n/n)^2 h(X_0)/h(X_n) \Rightarrow \beta^2 h(X_0)/h(X_\infty)$$

where X_∞ is a r.v. with mass function π_R and X_0 and X_∞ are independent. Since $|h(X_0)/h(X_n)| \leq \max\{h(x)/h(y) : x, y \in S\} < \infty$, we can apply the bounded convergence theorem to (3.4), yielding

$$(3.5) \quad n^{-2} E_R[Y_n^2 h(X_0)/h(X_n)] \rightarrow \beta^2 \kappa_1 \kappa_2,$$

where $\kappa_1 = \sum_x \eta(x)h(x)$ and $\kappa_2 = \sum_x \pi_R(x)/h(x)$. If $\beta \neq 0$, (3.5) provides the asymptotic estimate that we need (since the strict positivity of h implies that $\kappa_1, \kappa_2 > 0$). If $\beta = 0$, we note that (3.3) proves that under P_R ,

$$(Y_n/n^{\frac{1}{2}})^2 h(X_0)/h(X_n) \Rightarrow \sigma_R^2 N(0, 1)^2 h(X_0)/h(X_\infty),$$

where $N(0, 1)$, X_0 , and X_∞ are independent r.v.'s. Theorem 3, p. 102, of Chung (1967), together with the boundedness of $h(X_0)/h(X_n)$, then yields the conclusion that if $\beta = 0$,

$$n^{-1} E_R[Y_n^2 h(X_0)/h(X_n)] \rightarrow \sigma_R^2 \kappa_1 \kappa_2$$

as $n \rightarrow \infty$. We can summarize our discussion thus far with the following theorem.

THEOREM 2. Suppose P is irreducible and aperiodic. Then,

i) if $\beta \neq 0$ and $P \neq K$,

$$\text{var}_K[Y_n L_n] \sim n^2 \gamma \lambda^n \beta^2 \kappa_1 \kappa_2$$

as $n \rightarrow \infty$;

ii) if $\beta = 0$, $P \neq K$, and $\sigma_R^2 > 0$,

$$\text{var}_K[Y_n L_n] \sim n \gamma \lambda^n \sigma_R^2 \kappa_1 \kappa_2,$$

as $n \rightarrow \infty$.

Typically, we would expect $\beta \neq 0$ to hold in most practical settings, in which case the variance grows as $n^2 \lambda^n$.

It is instructive to also describe the asymptotic distribution of the r.v. $Y_n L_n$. Let $P_K(\cdot) = E_K[I(\cdot)]$ be the probability distribution on X associated with initial distribution ν and transition matrix K . Then,

$$\log(L_n) = \log \left[\frac{\nu(X_0)}{\mu(X_0)} \right] + \sum_{i=0}^{n-1} \log \left[\frac{P(X_i, X_{i+1})}{K(X_i, X_{i+1})} \right].$$

Applying the law of large numbers to the finite chain (X_n, X_{n+1}) , we conclude that

$$\frac{1}{n} \log L_n \rightarrow \sum_{x,y} \pi_K(x) K(x,y) \log \left[\frac{P(x,y)}{K(x,y)} \right] \triangleq \xi$$

P_K a.s. as $n \rightarrow \infty$. If $P \neq K$, the strict concavity of the log function implies that

$$\xi < \log \left[\sum_{x,y} \pi_K(x) k(x,y) \frac{P(x,y)}{K(x,y)} \right] = \log 1 = 0,$$

so that in this case,

$$(3.6) \quad L_n^{\frac{1}{n}} \rightarrow \exp(\xi) < 1.$$

P_K a.s. as $n \rightarrow \infty$. For an arbitrary function g , let $\|g\| = \sup\{|g(x)| : x \in S\}$. Since $|Y_n| \leq n \cdot \|f\|$, it is evident that $|Y_n L_n|^{\frac{1}{n}} \leq n^{\frac{1}{n}} \|f\|^{\frac{1}{n}} L_n^{\frac{1}{n}}$, so that (3.6) yields

$$(3.7) \quad \overline{\lim}_n |Y_n L_n|^{\frac{1}{n}} < 1$$

P_K a.s. Fix $0 < \varepsilon < 1$. If $\overline{\lim}_n |Y_n L_n| \geq \varepsilon$, then

$$\overline{\lim}_n |Y_n L_n|^{\frac{1}{n}} \geq \overline{\lim}_n \varepsilon^{\frac{1}{n}} = 1,$$

contradicting (3.7). Thus, (3.7) implies that $\overline{\lim}_n |Y_n L_n| = 0$, so that

$$(3.8) \quad Y_n L_n \rightarrow 0$$

P_K a.s. as $n \rightarrow \infty$, whenever $P \neq K$.

This gives us a more complete description of the asymptotic behavior of the r.v. $Y_n L_n$. While (3.7) and (3.8) state that $Y_n L_n$ is very small (with high probability) when n is large. Theorem 2 asserts that when $Y_n L_n$ is large (with small probability), it must be extremely large (in order that the variance grow geometrically fast). Thus, for large n , $Y_n L_n$ is a r.v. that takes on extremely large values with very small probability and small values with very high probability.

4. APPLICATION 2: TERMINAL COSTS

As in Section 3, let f be a real-valued function defined on the state space S of the Markov chain X . In this section, we are concerned with the terminal cost corresponding to the function f , namely

$$Y_n = f(X_n)$$

Assume P is irreducible and aperiodic. We may then apply Theorem 1 to obtain the identity

$$(4.1) \quad \text{var}_K[Y_n L_n] = \gamma \lambda^n E_R[f^2(X_n) h(X_0)/h(X_n)] - \alpha_n^2,$$

where $\alpha_n = E_p Y_n$. Since R is aperiodic (by virtue of the aperiodicity of P), it is evident that for each $x \in S$,

$$(4.2) \quad P_R\{X_n = y | X_0 = x\} \rightarrow \pi_R(y)$$

as $n \rightarrow \infty$. Relation (4.2) implies that

$$(4.3) \quad P_R\{X_0 = x, X_n = y\} \rightarrow \eta(x)\pi_R(y)$$

as $n \rightarrow \infty$. Thus, X_0 and X_n are asymptotically independent r.v.'s. Applying the continuous mapping principle to (4.3), we conclude that

$$(4.4) \quad f^2(X_n)h(X_0)/h(X_n) \Rightarrow f^2(X_\infty)h(X_0)/h(X_\infty)$$

as $n \rightarrow \infty$ (in P_R distribution), where X_0 and X_∞ are independent r.v.'s (with X_∞ having distribution π_R). The bounded convergence theorem then implies that

$$E_R[f^2(X_n)h(X_0)/h(X_n)] \rightarrow \kappa_1 \kappa_3$$

as $n \rightarrow \infty$, where $\kappa_1 = \sum_x \eta(x)h(x)$ and $\kappa_3 = \sum_x \pi_R(x)f^2(x)/h(x)$. This discussion has yielded the following theorem. (Note that α_n is bounded.)

THEOREM 3. Suppose P is irreducible and aperiodic. If $f \neq 0$ and $K \neq P$, then

$$\text{var}_K[f(X_n)L_n] \sim \gamma \lambda^n \kappa_1 \kappa_3$$

as $n \rightarrow \infty$.

The argument employed in Section 3 to study the "almost sure" behavior of $Y_n L_n$ is equally applicable here. In particular, if $K \neq P$, we may conclude that

$$f(X_n)L_n \rightarrow 0$$

P_K a.s. as $n \rightarrow \infty$. As in the case of cumulative costs, it therefore follows that when the time horizon n is large, $f(X_n)L_n$ takes on extremely large values with small probability and very small values with high probability.

5. APPLICATION 3: STEADY-STATE COSTS

In this section, we apply the results of Section 3 (on cumulative costs) to the analysis of importance sampling for steady-state costs. Given a real-valued function f defined on the state space of X , let

$$Y_n = \sum_{k=0}^{n-1} f(X_k)$$

be the cumulative cost corresponding to the function f . Let $P_P(\cdot) = E_P I(\cdot)$ be the probability distribution on X associated with initial distribution μ and transition matrix P . Assume that P is irreducible and aperiodic. Then, the law of large numbers applies to Y_n and we may assert that

$$(5.1) \quad Y_n/n \rightarrow \sum_x \pi_P(x) f(x) \triangleq r$$

P_P a.s. as $n \rightarrow \infty$, where $\pi_P(\cdot)$ is the (unique) stationary distribution of P . The constant r may therefore be interpreted as the steady-state cost associated with the performance measure f . We note that the bounded convergence theorem applies to (5.1), yielding

$$E_P[Y_n/n] \rightarrow r$$

as $n \rightarrow \infty$. Since $E_P Y_n = E_K L_n Y_n$, it follows that

$$E_K[Y_n L_n/n] \rightarrow r$$

as $n \rightarrow \infty$. Hence, the r.v. $Y_n L_n/n$ (when generated under P_K) can be used as an estimator for the steady-state mean r .

In particular, suppose that T represents the computer budget available to estimate the steady-state mean r . To simplify our analysis, we assume that exactly one transition of the chain X is generated per unit time. Given the budget T , we can generate (under P_K) $m = m(T)$ replicates of the chain X , each of length $n(T)$, where $n(T) = \lfloor T/m(T) \rfloor$. This results in the estimator

$$r(T) = \frac{1}{T} \sum_{i=1}^{m(T)} Y_{in(T)} L_{in(T)},$$

where $Y_{in(T)}L_{in(T)}$ is the i 'th independent replicate of the r.v. $Y_{n(T)}L_{n(T)}$. The mean square error (MSE) of $r(T)$ is given by

$$(5.2) \quad MSE_K[r(T)] = \text{var}_K[r(T)] + \text{bias}_K[r(T)]^2,$$

where $\text{bias}_K[r(T)] = E_K r(T) - r$. We note that

$$\text{var}_K[r(T)] = \frac{m(T)}{T^2} \text{var}_K[Y_{n(T)}L_{n(T)}]$$

$$\text{bias}_K[r(T)] = E_P[Y_{n(T)}/n(T)] - r.$$

In our subsequent analysis of (5.2), we shall assume that $n(T) \rightarrow \infty$ as $T \rightarrow \infty$. (In virtually all practical applications, this is necessary in order that the MSE converge to zero as $T \rightarrow \infty$.)

Now, Theorem 2 states that if $P \neq K$ and $\beta \neq 0$, then

$$(5.3) \quad \text{var}_K[r(T)] \sim \gamma \beta^2 \kappa_1 \kappa_2 n(T) \lambda^{n(T)} / T.$$

To analyze the bias term, we recall that we are assuming that P is aperiodic and irreducible. Then, it is well known that $E_P f(X_n) \rightarrow r$ geometrically fast, and hence

$$b = \sum_{n=0}^{\infty} (E_P f(X_n) - r)$$

converges absolutely. Hence, there exists $\rho \in [0, 1)$ such that

$$\begin{aligned} E_P[Y_n/n] - r &= b/n - \sum_{j=n}^{\infty} (E_P f(X_j) - r)/n \\ &= b/n + O(\rho^n) \end{aligned}$$

as $n \rightarrow \infty$. Thus, if $b \neq 0$ (as is typical in most applications), it is evident that

$$(5.4) \quad (\text{bias}_K[r(T)])^2 \sim b^2/n^2(T)$$

as $T \rightarrow \infty$. The following theorem is easily verified from the asymptotic formulae (5.3) and (5.4).

THEOREM 4. Suppose that P is irreducible and aperiodic. Assume $b \neq 0$, $\beta \neq 0$, $P \neq K$, and that $n(T) \rightarrow +\infty$.

i) If $n(T) - (\log T - \log \log T)/\log \lambda \rightarrow +\infty$ as $T \rightarrow \infty$, then

$$\text{var}_K[r(T)] \rightarrow +\infty, \text{bias}_K[r(T)] \rightarrow 0$$

as $T \rightarrow \infty$.

ii) If $n(T) - (\log T - 3 \log \log T)/\log \lambda \rightarrow +\infty$ as $T \rightarrow \infty$, then

$$(\log T)^2 \text{var}_K[r(T)] \rightarrow \infty, \overline{\lim} (\log T)^2 \text{bias}_K^2[r(T)] \leq (\log \lambda)^2 b^2$$

as $T \rightarrow \infty$.

iii) If $n(T) - (\log T - 3 \log \log T)/\log \lambda \rightarrow a$ as $T \rightarrow \infty$, then

$$(\log T)^2 \text{var}_K[r(T)] \rightarrow \gamma \beta^2 \kappa_1 \kappa_2 e^a, (\log T)^2 \text{bias}_K^2[r(T)] \rightarrow (\log \lambda)^2 b^2$$

as $T \rightarrow \infty$.

iv) If $n(T) - (\log T - 3 \log \log T)/\log \lambda \rightarrow -\infty$ as $T \rightarrow \infty$, then

$$(\log T)^2 \text{var}_K[r(T)] \rightarrow 0, \underline{\lim} (\log T)^2 \text{bias}_K^2[r(T)] \geq (\log \lambda)^2 b^2$$

as $T \rightarrow \infty$.

v) If $n(T) \sim c \log T$, where $0 < c < 1/\log \lambda$, then

$$(\log T)^2 \text{var}_K[r(T)] \rightarrow 0, (\log T)^2 \text{bias}_K^2[r(T)] \rightarrow b^2/c^2$$

as $T \rightarrow \infty$.

This result has a number of important implications. Firstly, we note that part i) states that if $n(T) - (\log T - \log \log T)/\log \lambda \rightarrow \infty$, $MSE_K r(T) \rightarrow \infty$, so that if the replication run-length $n(T)$ grows too quickly, the mean square error need not even converge to zero. This is particularly problematic in light of the fact that $\log \lambda$ is typically unknown.

Parts ii) through iv) suggest that $n(T) \approx (\log T - 3 \log \log T)/\log \lambda + a$ defines the critical case in which both components of the mean square error (namely variance and the squares bias) go to zero at the same rate, namely $(\log T)^{-2}$.

Finally, the theorem (taken as a whole) suggests that the best possible convergence rate for the root mean square error of $r(T)$ is $(\log T)^{-1}$ as $T \rightarrow \infty$. Thus, the convergence

rate of the replicated steady-state importance sampling estimator is exceptionally slow (when compared to the rate of $T^{-1/2}$ which is achieved for typical steady-state simulations when implemented without importance sampling).

The poor convergence rate of $r(T)$ basically arises because of the exponential growth in the variance of $Y_n L_n$. One way to (partially) avoid this is to use the regenerative structure that is present in finite-state Markov chains. This allows us to reduce the time horizon to that of a regenerative cycle. Specifically, select a regeneration state x and let $\tau = \inf\{n \geq 1 : X_n = x\}$. Thus, if we set $\mu = \delta_x$ (i.e. a unit point mass at x), we recall that the steady-state cost r can be represented as

$$(5.5) \quad \begin{aligned} r &= E_P Y_\tau / E_P \tau \\ &= E_K[Y_\tau L_\tau] / E_P[\tau L_\tau]. \end{aligned}$$

As suggested in Glynn and Iglehart (1989), one can use importance sampling to estimate r via the ratio formula (5.5). Let $(Y_{1\tau}, \tau_1, L_{1\tau}), (Y_{2\tau}, \tau_2, L_{2\tau}), \dots$ be a sequence of i.i.d. replicates of (Y_τ, τ, L_τ) (generated under P_K). Suppose that the time required to generate the i 'th cycle is τ_i and let $N(T) = \inf\{n \geq 0 : \sum_{i=1}^n \tau_i \leq T\}$ be the number of cycles completed in T units of computer time. Let $r'(T)$ be the ratio estimator available after T units of computer time have been expended, namely

$$r'(T) = \begin{cases} \frac{\sum_{i=1}^{N(T)} Y_{i\tau} L_{i\tau}}{N(T)} & ; N(T) \geq 1 \\ \sum_{i=1} \tau_i L_{i\tau} & \\ 0 & ; N(T) = 0. \end{cases}$$

In Glynn and Iglehart (1989), it is shown that if $E_K Z^2 < \infty$, where $Z = (Y_\tau - r\tau)L_\tau$, then

$$T^{1/2}(r'(T) - r) \Rightarrow dN(0, 1)$$

as $T \rightarrow \infty$, where $d^2 = E_K \tau \cdot E_K Z^2 / (E_P \tau)^2$. Thus, the regenerative estimator $r'(T)$ enjoys a convergence rate of $T^{-1/2}$, provided that $E_K Z^2 < \infty$. Unfortunately, as argued in Glynn and Iglehart (1989), the quantity $E_K Z^2$ can be infinite (even in the current finite state Markov chain setting).

We can use the machinery developed in this paper to obtain a precise analysis of when $E_K Z^2 < \infty$ will hold. Let $f_c(x) = f(x) - r$ and observe that $Z = Y'_\tau L_\tau$, where

$$Y'_n = \sum_{k=0}^{n-1} f_c(X_k).$$

Then, for each $n \geq 1$, Theorem 1 applies, yielding

$$(5.6) \quad E_K[Y_n'^2 L_n^2 I(\tau = n)] = \gamma \lambda^n E_K \left[Y_n'^2 \frac{h(X_0)}{h(X_n)} I(\tau = n) \right].$$

Summing each side of (5.6) over n and using Fubini's theorem, we get the identity

$$(5.7) \quad E_K Z^2 = \gamma E_R \left[\lambda^\tau \left(\sum_{k=0}^{\tau-1} f_c(X_k) \right)^2 \right].$$

(We've used the fact that $h(X_0) = h(X_\tau)$ and that $P_K\{\tau < \infty\} = P_R\{\tau < \infty\} = 1$ in obtaining (5.7)). In virtually all practical applications, it will therefore be necessary that $E_R \lambda^\tau < \infty$, in order that $E_K Z^2 < \infty$. Furthermore, if $E_R \lambda_0^\tau < \infty$ for some $\lambda_0 > 0$, this will always be sufficient to guarantee the finiteness of $E_K Z^2$. (To see this, note that $\lambda^\tau Y_\tau'^2 \leq \|f_c\|^2 \tau^2 \lambda^\tau \leq \|f_c\|^2 \lambda_0^\tau$ on $\{\tau \geq n_0\}$, where n_0 is chosen sufficiently large.) Thus, the finiteness of $E_K Z^2$ basically comes down to the issue of when the Perron-Frobenius eigenvalue $\lambda = \lambda(G)$ lies in the interior of the set $\{z \in \mathbb{R} : E_R z^\tau < \infty\}$.

Let A be the submatrix of R defined by $A = (R(u, v) : u, v \in S - \{x\})$. (Note that A has one fewer row and column than does R .) Clearly, A is non-negative. Suppose that A is irreducible. Then, A has a unique positive Perron-Frobenius eigenvalue $\lambda(A)$.

PROPOSITION 2. Suppose that A is irreducible and finite. Then, $E_R \lambda^\tau < \infty$ if and only if $\lambda(G) < \lambda(A)^{-1}$.

For the proof, see the Appendix.

Thus, the slow convergence rate of the replicated steady-state estimator $r(T)$ manifests itself in the regenerative setting through the possibility of infinite variance. If $E_K Z^2 = +\infty$, one can not expect a convergence rate of $T^{-\frac{1}{2}}$, but must instead expect a slower rate of convergence. As we have just argued, in order that $E_K Z^2 < \infty$, this will typically require that $\lambda(G) < \lambda(A)^{-1}$. Unfortunately, since $\lambda(G)$ and $\lambda(A)$ are unknown in practical

settings, this suggests that a great deal of care must be exercised in applying importance sampling using the regenerative estimator $r'(T)$.

6. APPLICATION 4: LIKELIHOOD RATIO GRADIENT ESTIMATOR

Consider a family $\{(P(\theta), \mu(\theta)) : \theta \in \Lambda\}$ of transition matrices and initial distributions on S that are indexed by some open set $\Lambda \subseteq \mathbb{R}^k$. As mentioned in the Introduction, the calculation of the gradient of the expected performance of the Markov chain with respect to θ is a problem that has recently attracted considerable attention within the simulation community.

One approach to this problem is known as likelihood ratio gradient estimation. In order to simplify the notation, we specialize (without any essential loss of generality) to the scalar case in which $k = 1$. Given a performance measure $Y = f(X_0, X_1, \dots, X_n)$, its expected value under initial distribution $\mu(\theta)$ and transition matrix $P(\theta)$ is given by

$$\alpha(\theta) \triangleq \sum_{x_0, \dots, x_n} f(x_0, \dots, x_n) \mu(\theta, x_0) \prod_{i=0}^{n-1} P(\theta, x_i, x_{i+1}).$$

Assume that $P(\theta)$ and $\mu(\theta)$ are both continuously differentiable on Λ . Then, for $\theta_0 \in \Lambda$, $\alpha'(\theta_0)$ exists and is given by

$$(6.1) \quad \alpha'(\theta_0) = \sum_{x_0, \dots, x_n} f(x_0, \dots, x_n) \left[\mu'(\theta_0, x_0) \prod_{i=0}^{n-1} P(\theta_0, x_i, x_{i+1}) + \mu(\theta_0, x_0) \sum_{i=0}^{n-1} P'(\theta_0, x_i, x_{i+1}) \prod_{j \neq i} P(\theta_0, x_j, x_{j+1}) \right]$$

In order to estimate $\alpha'(\theta_0)$ via simulation, it is necessary to represent $\alpha'(\theta_0)$ as an expectation. Suppose that we select ν such that $\nu(x) > 0$ whenever $\mu'(\theta_0, x) \neq 0$ or $\mu(\theta_0, x) > 0$ and select K so that $K(x, y) \neq 0$ whenever $P'(\theta_0, x, y) \neq 0$ or $P(\theta_0, x, y) > 0$. An important observation here is that, under our hypotheses, selecting $\nu = \mu(\theta_0)$ and $K = P(\theta_0)$ always fits this prescription. (The key point is that if $P'(\theta_0, x, y) \neq 0$ when $P(\theta_0, x, y) = 0$, this implies that $P(\cdot, x, y)$ is strictly negative in some neighborhood of θ_0 .) Let $E_\theta(\cdot)$ be the expectation operator associated with initial distribution $\mu(\theta)$ and transition matrix $P(\theta)$. Then, (6.1) can be re-written as

$$(6.2) \quad \alpha'(\theta_0) = E_{\theta_0}[Y L'_n(\theta_0)]$$

where

$$L'_n(\theta_0) = \frac{\mu'(\theta_0, X_0)}{\mu(\theta_0, X_0)} + \sum_{i=0}^{n-1} \frac{P'(\theta_0, X_i, X_{i+1})}{P(\theta_0, X_i, X_{i+1})}.$$

The r.v. $L'_n(\theta_0)$ is known as the score function. We can now apply importance sampling (see formula (1.1)) to (6.2), thereby yielding

$$(6.3) \quad \alpha'(\theta_0) = E_K[Y L'_n(\theta_0) L_n],$$

where

$$L_n = \frac{\mu(\theta_0, X_0)}{\nu(X_0)} \prod_{i=0}^{n-1} \frac{P(\theta_0, X_i, X_{i+1})}{K(X_i, X_{i+1})}.$$

The estimator based on $Y L'_n(\theta_0) L_n$ (when generated under P_K) is called the likelihood ratio gradient estimator.

Suppose $P(\theta_0)$ is irreducible and finite. By Theorem 1, we arrive at the identity

$$(6.4) \quad \text{var}_K[Y L'_n(\theta_0) L_n] = \gamma \lambda^n E_R[Y^2 L'_n(\theta_0)^2 h(X_0)/h(X_n)] - \alpha'(\theta_0)^2$$

Since $\lambda > 1$ if $P(\theta_0) \neq K$, (6.4) strongly suggests that the choice $K = P(\theta_0)$ minimizes the variance of the likelihood ratio gradient estimator when the time horizon is large.

To obtain a more precise statement, we need to specify the performance measure. Specifically, let us consider a cumulative cost of the form

$$Y_n = \sum_{k=0}^{n-1} f(X_k),$$

where f is some real-valued function defined on the state space S of X . To analyze the right-hand side of (6.4), we observe that the finiteness and irreducibility of R guarantees that the following strong laws will hold:

$$(6.5) \quad \begin{aligned} \frac{1}{n} Y_n &\rightarrow \sum_x \pi_R(x) f(x) \triangleq \beta \\ \frac{1}{n} L'_n(\theta_0) &\rightarrow \sum_x \pi_R(x) R(x, y) \frac{P'(\theta_0, x, y)}{P(\theta_0, x, y)} \triangleq \psi \end{aligned}$$

P_R a.s. as $n \rightarrow \infty$. As in Section 3, one can show that X_0, X_n , and $(Y_n, L'_n(\theta_0))$ are asymptotically independent of one another. If $P(\theta_0)$ is additionally assumed to be aperiodic, then (6.5) and the continuous mapping principle implies that

$$n^{-4} Y_n^2 L'_n(\theta_0)^2 h(X_0)/h(X_n) \Rightarrow \beta^2 \psi^2 h(X_0)/h(X_\infty),$$

as $n \rightarrow \infty$, where X_∞ is a r.v. having mass function $\pi_R(\cdot)$. The bounded convergence theorem then yields

$$n^{-4} E_R[Y_n^2 L'_n(\theta_0)^2 h(X_0)/h(X_n)] \rightarrow \beta^2 \psi^2 \kappa_1 \kappa_2$$

as $n \rightarrow \infty$. We can summarize our discussion thus far with the following theorem.

THEOREM 5. Suppose that $P(\theta_0)$ is irreducible and aperiodic. If $K \neq P(\theta_0)$, $\beta \neq 0$, and $\psi \neq 0$, then

$$\text{var}_K[Y_n L'_n(\theta_0) L_n] \sim \gamma \lambda^n n^4 \beta^2 \psi^2 \kappa_1 \kappa_2$$

as $n \rightarrow \infty$.

In L'Ecuyer and Glynn (1991), it is shown that $\text{var}_{\theta_0}[Y_n L'_n(\theta_0)]$ is typically of order n^3 . Hence, Theorem 5 shows that choosing $K \neq P(\theta_0)$ significantly degrades the variance (for cumulative cost performance measures) when the time horizon is large.

We conclude this section with a brief discussion of the implications for optimization. Assume that $P(\cdot)$ is such that if $P(\theta, x, y) > 0$ for some θ , then $P(\theta, x, y) > 0$ for all θ . Similarly, assume that if $\mu(\theta, x) > 0$ for some θ , then $\mu(\theta, x) > 0$ for all θ . Suppose that, as suggested in the Introduction, one simulates the Markov chain X under the distribution P_{θ_0} associated with parameter point $\theta_0 \in \Lambda$. Then, we can obtain a global estimate for $\alpha'(\cdot)$ by using (6.3):

$$(6.6) \quad \alpha'(\theta) = E_{\theta_0}[Y L'_n(\theta) L_n(\theta, \theta_0)]$$

where

$$L'_n(\theta) = \frac{\mu'(\theta, X_0)}{\mu(\theta, X_0)} + \sum_{i=0}^{n-1} \frac{P'(\theta, X_i, X_{i+1})}{P(\theta, X_i, X_{i+1})}$$

$$L_n(\theta, \theta_0) = \frac{\mu(\theta, X_0)}{\mu(\theta_0, X_0)} \prod_{i=0}^{n-1} \frac{P(\theta, X_i, X_{i+1})}{P(\theta_0, X_i, X_{i+1})}.$$

Note that by simulating at the single parameter point θ_0 , we can obtain an unbiased estimate for $\alpha'(\theta)$ at each $\theta \in \Lambda$. However, as shown earlier, we can expect the variance

of the estimator for $\alpha'(\theta)$ to increase geometrically (in the length of the time horizon) at rate $\lambda(\theta)$, where $\lambda(\theta)$ is the Perron-Frobenius eigenvalue of the matrix $G(\theta) = (G(\theta, x, y) : x, y \in S)$ and $G(\theta, x, y) = P^2(\theta, x, y)/P(\theta_0, x, y)$. Recall that

$$\sum_y G(\theta, x, y) - 1 = \sum_y \left[\frac{P(\theta, x, y)}{P(\theta_0, x, y)} - 1 \right]^2 P(\theta_0, x, y).$$

Thus, we can expect that as the distance between θ and θ_0 grows, the row sums of $G(\theta)$ grow. Let $\delta(\theta) = \min \left\{ \sum_y G(\theta, x, y) : x \in S \right\}$ and note that

$$G(\theta)e \geq \delta(\theta)e$$

where e is a column vector of 1's. Perron-Frobenius theory implies that if $P(\theta_0)$ is irreducible, there exists a positive row vector $x(\theta)$ such that $x(\theta)G(\theta) = \lambda(\theta)x(\theta)$. Hence,

$$\lambda(\theta)x(\theta)e = x(\theta)G(\theta)e \geq \delta(\theta)x(\theta)e.$$

Since $x(\theta)e > 0$, it follows that $\lambda(\theta) \geq \delta(\theta)$. Thus, the growth of the row sums forces $\lambda(\theta)$ to grow. So, we can expect that in many practical settings, $\lambda(\theta)$ will be large at points θ that are distant from θ_0 . Thus, the geometric growth problem discussed above may be particularly troublesome at points θ that are distant from θ_0 . This suggests that great care needs to be taken with the development of optimization algorithms based on (6.6).

7. IMPLICATIONS FOR DISCRETE-EVENT SIMULATION

Thus far, our discussion in this paper has focused on the analysis of importance sampling, as it applies to discrete-time finite state Markov chains. However, we believe that the results presented here have obvious analogues in the more general discrete-event simulation context.

The basic idea is that a typical discrete-event simulation, when considered on the time scale of state transitions, can be viewed as a discrete time Markov chain living on a general (uncountable) state space. In particular, suppose that we let S_n represent the "physical state" (e.g., the location of the customers in a queuing network, as described by a queue-length vector) of the system at the time of the n 'th (physical) state transition. Also, let

C_n be the “clock reading” vector, at the time of the n 'th transition. The i th component of C_n then describes the time that remains until the i 'th possible trigger event will initiate a state transition. The key observation is that $X_n = (S_n, C_n)$ is then a discrete-time Markov chain; see Glynn (1983) and Glynn (1989) for further details. Of course, the state space of this Markov chain is typically uncountable.

The analysis presented in this paper hinges on two important pieces of mathematical machinery. The first is the existence of the Perron-Frobenius theory of non-negative matrices. Fortunately, much of this theory carries over to the more general setting of non-negative operators acting on an abstract state space; see Chapter 5 of Nummelin (1984) for a recent account of this theory. The second tool that was repeatedly applied was the limit theory for additive functionals of finite state Markov chain (e.g., laws of large numbers and central limit theorems). Again, these results have a number of generalizations to the uncountable state space setting; see Niemi and Nummelin (1982) for a description of such results.

Our view is therefore that, under suitable regularity hypotheses, the results of this paper will carry over to the general state space setting, and hence to discrete-event simulations. For example, we would expect (as in Theorem 2) that when importance sampling is applied to a cost that is cumulated over the first n transitions of the system, the variance will typically grow at rate $n^2 \lambda^n$ for some constant $\lambda > 1$. In other words, the results obtained in this paper are qualitatively representative of what one should expect in the more general discrete-event simulation setting.

APPENDIX

Proof of Proposition 1. Consider a typical row of the matrix G . Then, for each $x \in S$,

$$(A.1) \quad \begin{aligned} \sum_y G(x, y) - 1 &= \sum_y P^2(x, y)/K(x, y) - 1 \\ &= \sum_y \left[\frac{P(x, y)}{K(x, y)} - 1 \right]^2 K(x, y) \geq 0. \end{aligned}$$

Let $B = (B(x, y) : x, y \in S)$ be the stochastic matrix defined by

$$B(x, y) = G(x, y) / \sum_z G(x, z),$$

and note that $G(x, y) \geq B(x, y)$ for each $x, y \in S$ (since the normalization factor that defines each row of B is greater than or equal to 1, by (A.1)). By Corollary 2.3, p. 551, of Karlin and Taylor (1975), $\lambda = \lambda(G) \geq \lambda(B)$, where $\lambda(B)$ is the Perron-Frobenius eigenvalue of B . But since B is stochastic, $\lambda(B) = 1$. Hence, $\lambda \geq 1$.

As for γ , an argument similar to (A.1) shows that $\gamma \geq 1$, with $\gamma = 1$ if and only if $\nu = \mu$.

It remains to show that $\lambda(G) > 1$ if $P \neq K$. In this case, at least one row sum of G is strictly greater than 1. Hence, there exists at least one state x_* such that $G(x_*, y) > B(x_*, y)$ for each state y such that $B(x_*, y) > 0$.

Note that B is irreducible, since P (and hence G) is. Since B is stochastic, there exists a unique strictly positive stochastic row vector φ (the stationary distribution of B) such that $\varphi B = \varphi$. Furthermore, there exists a strictly positive column eigenvector h corresponding to the Perron-Frobenius eigenvalue $\lambda = \lambda(G)$ of G . Since $Gh = \lambda h$, it is evident that

$$(A.2) \quad \varphi Gh = \lambda \varphi h.$$

On the other hand, $Gh \geq Bh$ with strict inequality in the x_* 'th component. (This follows from the strict positivity of h .) Hence, $\varphi Gh > \varphi Bh$. (Here, we make use of the strict positivity of φ .) So, using the stationarity of φ and (A.2), we get

$$\lambda \varphi h > \varphi Bh = \varphi h.$$

We conclude that $\lambda = \lambda(G) > 1$ (since the positivity of φ and h implies that $\varphi h > 0$).

Proof of Proposition 2. The argument is very similar to that used in Section 2 to study $\text{var}_K[YL_n]$. We note that

$$E_R \lambda^\tau = \lambda + \lambda \sum_{y \neq x} R(x, y) E_R[\lambda^\tau | X_0 = y].$$

Recall that A necessarily possesses a strictly positive column eigenvector h' such that $Ah' = \lambda(A)h'$. Set $C = (C(u, v) : u, v \in S - \{x\})$, where

$$C(u, v) = \frac{A(u, v)h'(v)}{\lambda(A)h'(u)}.$$

Arguing as in Section 2, it is easy to see that for $y \neq x$, $n \geq 1$,

$$P_R\{\tau = n | X_0 = y\} = \lambda(A)^{n-1} E_C \left[\frac{h'(X_0)}{h'(X_{n-1})} R(X_{n-1}, x) | X_0 = y \right].$$

where $E_C(\cdot)$ is the expectation operator in which X evolves according to the (stochastic) matrix C . We therefore arrive, for $y \neq x$, at the identity

$$(A.3) \quad \begin{aligned} E_R[\lambda^\tau | X_0 = y] &= \lambda(A)R(y, x) \\ &+ \lambda(A) \sum_{n=0}^{\infty} [\lambda(G)\lambda(A)]^n E_C \left[\frac{h'(X_0)}{h'(X_n)} R(X_n, x) | X_0 = y \right]. \end{aligned}$$

Note that the expectations $E_C(\cdot)$ appearing in (A.3) are bounded above by $\max\{h'(u)R(v, x)/h'(v) : u, v \in S - \{x\}\}$ so that if $\lambda(G)\lambda(A) < 1$, (A.3) clearly converges. On the other hand, there exists a sequence of the form $I = \{i + jm : m \geq 0\}$ (depending on the periodicity of X under C) such that $R(X_n, x)$ is bounded away from zero for $n \in I$. Since h' is strictly positive, this guarantees that the expectations $E_C(\cdot)$ are bounded away from zero on the subsequence I . Thus, in order that (A.3) converge, it is necessary that $\lambda(G)\lambda(A) < 1$.

REFERENCES

- Billingsley, P. (1968). *Convergence of Probability Measures*. John Wiley, New York.
- Bucklew, J. A. (1990). *Large Deviations Techniques in Decision, Simulation, and Estimation*. John Wiley, New York.
- Chung, K. L. (1967). *Markov Chains with Stationary Transition Probabilities*. Springer-Verlag, New York.
- Cottrell, M. J. C. Fort, and G. Malgouyres (1983). Large deviations and rare events in the study of stochastic algorithms. *IEEE Transactions on Automatic Control AC-28*, 907-920.
- Glynn, P. W. (1983). On the role of generalized semi-Markov processes in simulation output analysis. *Proc. of the 1983 Winter Simulation Conference*, 38-42.
- Glynn, P. W. (1986a). Sensitivity analysis for stationary probabilities of a Markov chain. *Proc. of the 4'th Army Conference on Applied Mathematics and Computing*, 917-932.
- Glynn, P. W. (1986b). Stochastic approximation for Monte Carlo optimization. *Proc. of the 1986 Winter Simulation Conference*, 356-365.
- Glynn, P. W. (1987). Likelihood ratio gradient estimation: An overview. *Proc. of the 1987 Winter Simulation Conference*, 318-319.
- Glynn, P. W. (1989). A GSMP formalism for discrete-event systems. *Proceedings of the IEEE*, 14-23.
- Glynn, P. W. and D. L. Iglehart (1989). Importance sampling for stochastic simulation. *Management Science* **35**, 1367-1392.
- Goyal, A., P. Shahabuddin, P. Heidelberger, V. F. Nicola, and P. W. Glynn (1991). A unified framework for simulating Markovian models of highly dependable systems. To appear in *IEEE Transactions on Computers*.
- Karlin, S. and H. J. Taylor (1975). *A First Course in Stochastic Processes*. Academic Press, New York.
- L'Ecuyer, P. and P. W. Glynn (1991). A control variate scheme for likelihood ratio gradient estimators. Manuscript in preparation.
- Niemi, S. and E. Nummelin (1982). Central limit theorems for Markov random walks. *Commentationes Physico-Mathematicae* **54**, Societas Scientiarum Fennica, Helsinki.

- Nummerlin, E. (1984). *General Irreducible Markov Chains and Non-Negative Operators*. Cambridge University Press, Cambridge.
- Rubinstein, R. Y. and A. Shapiro (1989). Optimization of simulation models by the score function method. Mimeograph Series 397, Faculty of Industrial Engineering and Management, Technion-Israel Institute of Technology, Israel.